

Exhibit D

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner,

v.

SINGULAR COMPUTING LLC,
Patent Owner.

IPR2021-00165
Patent 9,218,156 B2

Before JUSTIN T. ARBES, KRISTIL R. SAWERT, and
JASON M. REPKO, *Administrative Patent Judges*.

PER CURIAM.

DECISION
Granting Institution of *Inter Partes* Review
35 U.S.C. § 314

I. INTRODUCTION

Google LLC (“Petitioner”) filed a petition to institute *inter partes* review of claims 1–8, 16, and 33 of U.S. Patent No. 9,218,156 B2 (Ex. 1001, “the ’156 patent”). Paper 2 (“Pet.”). Singular Computing LLC (“Patent Owner”) filed a Preliminary Response. Paper 7 (“Prelim. Resp.”). On our authorization (Paper 13), Petitioner filed a Reply (Paper 14, “Reply”) and Patent Owner filed a Sur-Reply (Paper 15, “Sur-Reply”).

Under 35 U.S.C. § 314 and 37 C.F.R. § 42.4(a), we have authority to institute an *inter partes* review if “the information presented in the petition . . . and any response . . . shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” 35 U.S.C. § 314(a). After considering the Petition, the Preliminary Response, Reply, Sur-Reply, and the evidence of record, we determine the information presented shows a reasonable likelihood that Petitioner would prevail in establishing the unpatentability of at least one of the challenged claims of the ’156 patent. Accordingly, we institute an *inter partes* review of claims 1–8, 16, and 33 of the ’156 patent on the grounds asserted in the Petition.

II. BACKGROUND

A. Related Matters

The parties identify the following district-court proceeding as a related matter involving the ’156 patent: *Singular Computing LLC v. Google LLC*, No. 1:19-cv-12551-FDS (D. Mass.). Pet. xi; Paper 6, 1 (Patent Owner’s Mandatory Notices). The ’156 patent is also subject to a Petition for *inter partes* review in IPR2021-00164. Pet. x; Paper 6, 1. The following Board proceedings involve the same parties and related patents: IPR2021-

IPR2021-00165

Patent 9,218,156 B2

00154 and IPR2021-00155 (U.S. Patent No. 10,416,961 B2); and IPR2021-00178 and IPR2021-00179 (U.S. Patent No. 8,407,273 B2). Pet. x–xi; Paper 6, 1.

Petitioner also identifies as related matters the following patent applications, which claim priority to the same provisional application as the ’156 patent: U.S. Patent Application Nos. 14/976,852, 15/784,359, 16/175,131, 16/571,871, 16/675,693, 16/882,686, 16/882,694, and 17/029,780. Pet. x.

B. Overview of the ’156 Patent

The ’156 patent, entitled “Processing with Compact Arithmetic Processing Element,” relates to “computer processors or other devices which use low precision high dynamic range (LPHDR) processing elements to perform computations (such as arithmetic operations).” Ex. 1001, code [54], 6:3–7.

According to the ’156 patent, conventional CPU chips make inefficient use of transistors as a tradeoff for delivering the high precision required by many applications. *Id.* at 3:11–26. For example, conventional CPU chips “perform[] exact arithmetic with integers typically 32 or 64 bits long and perform[] rather accurate and widely standardized arithmetic with 32 and 64 bit floating point numbers,” but require “on the order of a million transistors to implement the arithmetic operations.” *Id.* at 3:19–26.

According to the ’156 patent, “many economically important applications . . . are not especially sensitive to precision and . . . would greatly benefit, in the form of application performance per transistor, from the ability to draw upon a far greater fraction of the computing power inherent in those million

IPR2021-00165

Patent 9,218,156 B2

transistors.” *Id.* at 3:27–32. But “[c]urrent architectures for general purpose computing fail to deliver this power.” *Id.* at 32–33.

The ’156 patent is therefore “directed to a processor or other device, such as a programmable and/or massively parallel processor or other device, which includes processing elements designed to perform arithmetic operations . . . on numerical values of low precision but high dynamic range (‘LPHDR arithmetic’).” *Id.* at 2:15–22. According to the ’156 patent, “‘low precision’ processing elements perform arithmetic operations which produce results that frequently differ from exact results by at least 0.1%.” *Id.* at 2:32–35. In addition, “high dynamic range” processing elements “are capable of operating on inputs and/or producing outputs spanning a range at least as large as from one millionth to one million.” *Id.* at 2:40–43. Figure 6, reproduced below, is an example of an LPHDR arithmetic unit according to one embodiment of the ’156 patent.

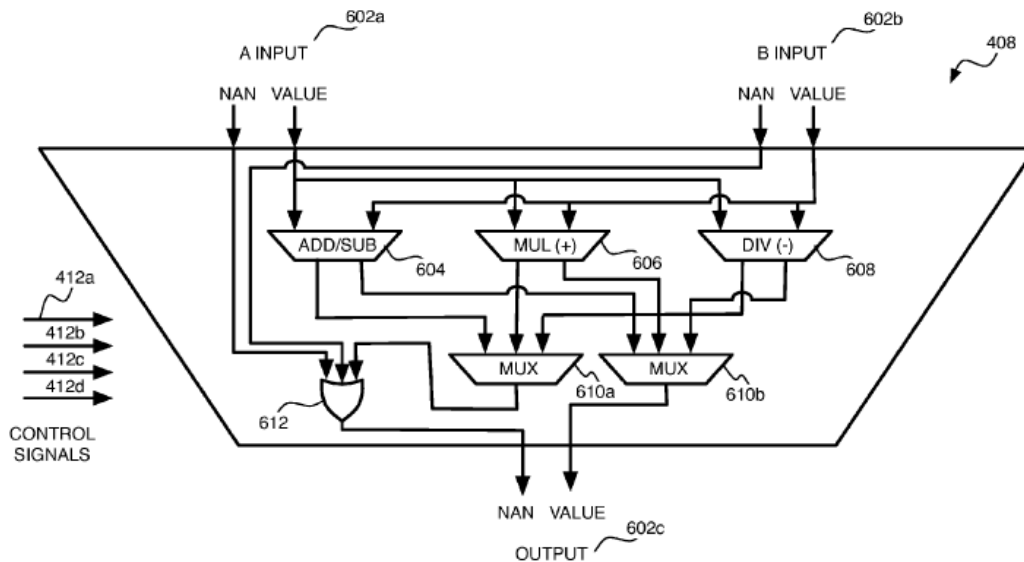


FIG. 6 provides “an example design for an LPHDR arithmetic unit according to one embodiment of” the ’156 patent. Ex. 1001, 2:60–61.

As shown in Figure 6, LPHDR arithmetic unit 408 receives two inputs: A input (602a) and B input (602b), and produces output 602c. *Id.* at 12:55–56. The LPHDR arithmetic unit “is controlled by control signals 412a-d, coming from the CU 106, that determine which available arithmetic operation will be performed on the inputs 602a-b.” *Id.* at 12:62–65. According to the ’156 patent, Figure 6 illustrates an embodiment where “all the available arithmetic operations are performed in parallel on the inputs 602a-b by adder/subtractor 604, multiplier 606, and divider 608.” *Id.* at 12:65–13:1. Finally, multiplexers (MUXes) 610a and 610b choose and send the desired result from among the outputs of the adder/subtractor, multiplier, and divider to output 602c. *Id.* at 13:4–13. The ’156 patent provides that “[t]he computing architecture literature discusses many variations which may be incorporated into the embodiment illustrated in FIG. 6.” *Id.* at 13:11–13.

According to the ’156 patent, the “computational tasks” that the LPHDR arithmetic units can perform “enable a variety of practical applications.” *Id.* at 17:20–23. The ’156 patent provides, as examples, applications including “finding nearest neighbors,” *id.* at 17:31–21:26, “distance weighted scoring,” *id.* at 21:28–22:17, and “removing motion blur in images,” *id.* at 22:19–23:34.

C. The Challenged Claims

Petitioner challenges claims 1–8, 16, and 33 of the ’156 patent. Pet. 4–6. Of the challenged claims, claims 1, 16, and 33 are independent. Claim 1, reproduced below, is illustrative of the claimed subject matter.

1. A device comprising:

at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input

IPR2021-00165

Patent 9,218,156 B2

signal representing a first numerical value to produce a first output signal representing a second numerical value,

wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from $1/65,000$ through $65,000$ and for at least $X=5\%$ of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least $X\%$ of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least $Y=0.05\%$ from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and

at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.

Ex. 1001, 29:54–30:6.

D. Evidence

Petitioner submits the following evidence:

Evidence	Exhibit No.
Declaration of Richard Goodin, P.E.	1003
U.S. Patent Appl. Pub. No. 2007/0203967 A1 (published Aug. 30, 2007) (“Dockser”)	1007
Tong et. al., <i>Reducing Power by Optimizing the Necessary Precision/Range of Floating-Point Arithmetic</i> , IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS, 8(3):273–286 (June 2000) (“Tong”)	1008
U.S. Patent No. 5,689,677 (issued Nov. 18, 1997) (“MacMillan”)	1009

IPR2021-00165

Patent 9,218,156 B2

E. Asserted Grounds of Unpatentability

Petitioner asserts the following grounds of unpatentability:

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
1, 2, 16	103(a) ¹	Dockser
1, 2, 16, 33	103(a)	Dockser, Tong
1–8, 16	103(a)	Dockser, MacMillan
1–8, 16, 33	103(a)	Dockser, Tong, MacMillan

Pet. 4–6. Patent Owner disputes Petitioner’s asserted grounds of unpatentability. *See generally* Prelim. Resp.

III. DISCUSSION

Petitioner contends that claims 1–8, 16, and 33 of the ’156 patent are unpatentable under 35 U.S.C. § 103 as obvious over various combinations of prior-art references Dockser, Tong, and MacMillan. A patent claim is unpatentable under § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and

¹ Congress amended §§ 102 and 103, among other sections, when it passed the Leahy-Smith America Invents Act (“AIA”). Pub. L. No. 112–29, § 3(c), 125 Stat. 284, 287 (2011). Here, Petitioner’s challenges are based on pre-AIA statutes. *See* Pet. 4.

(4) when in evidence, objective indicia of non-obviousness.² *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

“In an [*inter partes* review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.” *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016) (citing 35 U.S.C. § 312(a)(3) (requiring *inter partes* review petitions to identify “with particularity . . . the evidence that supports the grounds for the challenge to each claim”)). This burden of persuasion never shifts to Patent Owner. *See Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015) (discussing the burden of proof in *inter partes* review).

We organize our patentability analysis into four sections. First, we address the level of ordinary skill in the art. Second, we address claim construction. Third, we provide an overview of the asserted references. And fourth, taking account of the information presented, we consider whether the Petition satisfies the threshold requirement for instituting an *inter partes* review under 35 U.S.C. § 314(a).

A. Level of Ordinary Skill in the Art

We consider the asserted ground of unpatentability in view of the understanding of a person of ordinary skill in the art. The level of ordinary skill in the art is a factual determination that provides a primary guarantee of objectivity in an obviousness analysis. *Al-Site Corp. v. VSI Int’l Inc.*, 174

² With respect to the fourth *Graham* factor, the parties at this time do not present arguments or evidence regarding objective indicia of non-obviousness. Therefore, the obviousness analysis at this stage of the proceeding is based on the first three *Graham* factors.

F.3d 1308, 1324 (Fed. Cir. 1999) (citing *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966); *Ryko Mfg. Co. v. Nu-Star, Inc.*, 950 F.2d 714, 718 (Fed. Cir. 1991)).

Relying on the declaration testimony of Mr. Goodin, Petitioner contends that an ordinarily skilled artisan for the ’156 patent “would have had at least a bachelor’s degree in Electrical Engineering, Computer Engineering, Applied Mathematics, or the equivalent, and at least two years of academic or industry experience in computer architecture.” Pet. 8–9 (citing Ex. 1003 ¶¶ 43–44). Patent Owner does not address the level of ordinary skill in the art in its Preliminary Response. *See generally* Prelim. Resp.

Based on the record presented, including our review of the ’156 patent (including the types of problems and solutions described in the ’156 patent), and cited prior art, we agree with Petitioner’s proposed definition of the level of ordinary skill in the art with one exception. Arguably, the term “at least” creates unnecessary ambiguity. Thus, we delete that term from Petitioner’s definition, and otherwise apply Petitioner’s definition for this decision. *See, e.g.*, Ex. 1001, 1:31–2:10 (describing in the “Background” section of the ’156 patent various conventional methods of computation and their alleged deficiencies).

B. Claim Construction

In interpreting the claims of the ’156 patent, we “us[e] the same claim construction standard that would be used to construe the claim[s] in a civil action under 35 U.S.C. [§] 282(b).” 37 C.F.R. § 42.100(b) (2020). The claim construction standard includes construing claims in accordance with the ordinary and customary meaning of such claims as would have been

understood by one of ordinary skill in the art and the prosecution history pertaining to the patent. *See id.*; *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312–14 (Fed. Cir. 2005) (en banc).

Petitioner contends that the terms of the challenged claims should be “given their ordinary and customary meaning as understood by a [person of ordinary skill in the art] in accordance with the specification and prosecution history,” but does not propose any express interpretations. Pet. 9. Patent Owner disputes the construction of “low precision high dynamic range (LPHDR) execution unit” that Petitioner proposed in the related district court proceeding (i.e., “low precision and high dynamic range processing element designed to perform arithmetic operations on numerical values”). Prelim. Resp. 13–14 (citing Ex. 2001, 13–17 (emphasis omitted)). Patent Owner argues that Dockser does not teach an LPHDR execution unit because “the processing element itself must be fairly characterized as ‘low precision’” and cannot be an execution unit “whose subprecision can be selectively reduced.” *Id.* at 14 (emphasis omitted). We conclude that no terms require express interpretation at this time, and address the parties’ arguments regarding whether Dockser teaches the recited LPHDR execution unit below. *See infra* § III.D.

C. The Asserted Prior Art

Before turning to Petitioner’s asserted grounds of unpatentability, we provide a brief summary of the asserted references.

1. Dockser

Dockser discloses performing floating-point operations with a floating-point processor having selectable precision. Ex. 1007, code (57). Figure 1 of Dockser is reproduced below.

IPR2021-00165

Patent 9,218,156 B2

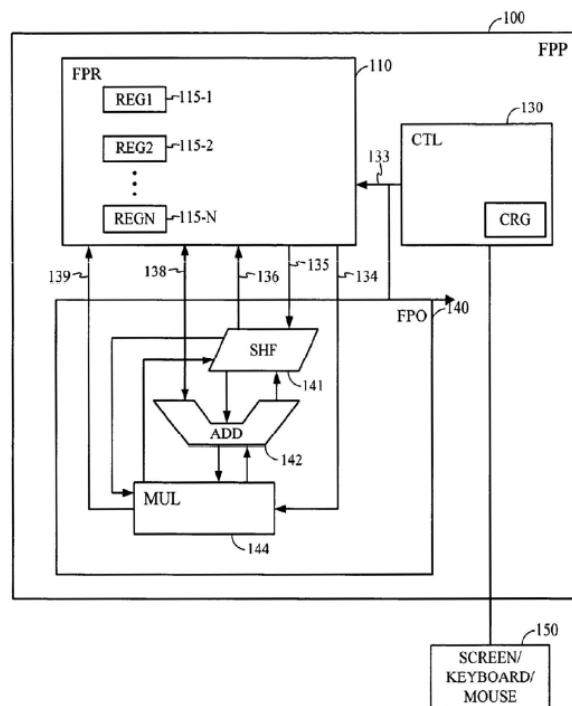


FIG. 1

Figure 1 depicts floating-point processor (FPP) 100 including floating-point register file (FPR) 110 for storing floating-point numbers, floating-point controller (CTL) 130 “used to select the subprecision of the floating-point operations using a control signal 133,” and floating-point mathematical operator (FPO) 140 with components “configured to perform the floating-point operations,” such as floating-point adder (ADD) 142 and floating-point multiplier (MUL) 144. *Id.* ¶¶ 15, 18, 19.

Figure 2 of Dockser is reproduced below.




FIG. 2

Figure 2 depicts an exemplary data structure for floating-point register file 110 including 16 addressable register locations 200, each “configured to store a 32-bit binary floating-point number” as “a 1-bit sign 202, an 8-bit exponent 204, and a [23-bit] fraction 206.” *Id.* ¶ 17.

“[F]or each instruction of a requested floating-point operation, the relevant computational unit . . . receive[s] from the floating-point register file 110 one or more operands stored in one or more of the register locations” and executes the instruction “at the subprecision selected by the floating-point controller 130.” *Id.* ¶¶ 23–24. The precision of the floating-point operation can be reduced by “caus[ing] power to be removed from the floating-point register elements for the excess bits of the fraction that are not required to meet the precision specified by the subprecision select bits” written to the control register. *Id.* ¶¶ 6, 25–26. For example, “if each location in the floating-point register file contains a 23-bit fraction, and the subprecision required for the floating-point operation is 10-bits, only the

9 commonly significant bits (MSBs) of the fraction are required; the hidden or integer bit makes the tenth.” *Id.* ¶ 26. “Power can be removed from the floating-point register elements for the remaining 14 fraction bits.” *Id.*

Alternatively, power can be removed in elements of “the logic in the floating-point operator 140 that remains unused as a result of the subprecision selected.” *Id.* ¶¶ 7, 27, 29, 32, *see also* Fig. 2 (depicting mantissa fraction 206 as having portion 322 for powered bits and portion 324 for unpowered bits). Figures 3A and 3B of Dockser show such removal of power to the floating-point operator logic for a floating-point addition and floating-point multiplication operation, respectively. Figure 3B of Dockser is reproduced below.




Figure 3A depicts k-bit multiplicand 402 and k-bit multiplier 404 to be multiplied together “using a shift-and-add technique,” where the multiplication occurs in stages 410-1 through 410-m. *Id.* ¶¶ 30–31. A partial product 420-i is generated for every bit in multiplier 404 at corresponding stage 410-i and then left-shifted “as a function of the multiplier bit with which it is associated, after which the operation moves on to the next stage.” *Id.* ¶ 31. The partial products are eventually added

together to generate output value 430. *Id.* “[P]ower may be removed from the logic used to implement the stages to the right of the line 405” indicating the selected subprecision. *Id.* ¶¶ 32–33.

2. *Tong*

Tong describes reducing power consumption by minimizing the bitwidth representation of floating-point data. Ex. 1008, 273. According to Tong, using a variable bitwidth floating-point unit saves power. *Id.*

3. *MacMillan*

MacMillan describes using Single Instruction Multiple Data (SIMD) parallel-processing architectures for adding supercomputer performance to personal-use computers. Ex. 1009, 5:22–54. MacMillan’s computer system comprises a “Host CPU” (i.e., “a 386, 486 or Pentium[] processor”) and SIMD-random access memory (SIMD-RAM) device. *Id.* at 9:30–31, Figs. 3, 5. MacMillan describes an example architecture where the SIMD-RAM device has 256 processing elements (PEs), but states that the disclosed architecture “allows scaling to higher or lower density chips with more or fewer PEs.” *Id.* at 12:60–13:4, 13:39–41, 16:20–22.

D. *Alleged Grounds of Unpatentability*

We now consider whether the Petition satisfies the threshold requirement for instituting an *inter partes* review under 35 U.S.C. § 314(a) by addressing each of Petitioner’s asserted grounds of unpatentability, below.

1. *Asserted Obviousness Over Dockser*

Petitioner contends that claims 1, 2, and 16 of the ’156 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over Dockser. Pet. 9–40. Patent Owner opposes. Prelim. Resp. 8–24. On the present record, for the

reasons detailed below, we determine that Petitioner has shown sufficiently for institution that Dockser teaches or suggests the limitations of the challenged claims.

a) Claim 1

Petitioner argues that Dockser teaches or suggests all the limitations of claim 1. Pet. 13–38. Claim 1 recites “[a] device comprising: at least one first low precision high dynamic range (LPHDR) execution unit.” Ex. 1001, 29:54–56. Petitioner contends that Dockser teaches a “device” (i.e., computing system) comprising a “low precision high dynamic range (LPHDR) execution unit” (i.e., the FPP).³ Pet. 13–15. With respect to “low precision,” Petitioner contends that the FPP is “low precision” because “‘the precision’ of operations in the FPP is ‘reduced’” and because the FPP “operates with the minimum imprecision” required by the subsequent language in claim 1. *Id.* at 14–15 (quoting Ex. 1007 ¶ 14). With respect to “high dynamic range,” Petitioner contends that the FPP “uses an 8-bit floating-point exponent . . . that provides an even higher dynamic range” than the 6-bit floating-point exponent disclosed in the ’156 patent. *Id.* at 15 (citing Ex. 1007 ¶ 17; Ex. 1001, 14:56–64).

Claim 1 recites that the LPHDR execution unit is “adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value.”

³ Petitioner also provides an “alternative mapping” where “the floating-point operator (FPO) inside Dockser’s FPP” constitutes an LPHDR execution unit. Pet. 14–15, 19, 21, 37. We need not evaluate those arguments at this time, because we determine that Petitioner has made a sufficient showing on the current record that the FPP is an LPHDR execution unit.

Ex. 1001, 29:56–59. Petitioner contends that Dockser’s FPP is adapted to execute a “first operation” (e.g., “reduced-precision multiplication”) on a “first input signal representing a first numerical value” (i.e., input electrical signal representing an operand received at the registers) to produce a “first output signal representing a second numerical value” (i.e., output electrical signal representing an operand sent to a register and then main memory), where the FPP “performs operations on the[] inputs via the FPP’s data paths 134–139 and components 140–144.” Pet. 15–19.

Next, claim 1 recites that the “dynamic range of the possible valid inputs to the first operation is at least as wide as from $1/65,000$ through $65,000$.” Ex. 1001, 29:60–62. Petitioner contends that Dockser’s “dynamic range of the possible valid inputs to the first operation is at least as wide as from $1/65,000$ through $65,000$,” as recited in claim 1, because the FPP “operates on IEEE-754 32-bit single-format numbers having 8-bit exponents” and, therefore, the dynamic range of normal operands would be “from around 2^{-126} (much smaller than $1/65,000$) to around 2^{127} (much larger than $65,000$).” Pet. 20.

The next limitation of claim 1 recites that

for at least $X=5\%$ of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least $X\%$ of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least $Y=0.05\%$ from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.

Ex. 1001, 29:62–30:4. We refer to this limitation as “the imprecision limitation.” With respect to the “statistical mean” portion of the imprecision

limitation, Petitioner contends that an ordinarily skilled artisan would have understood the limitation

in the context of the '156 patent's stated intent to claim not only "repeatable" deterministic embodiments like digital circuits that always produce the same output when repeating an operation on the same input, but also analog embodiments that are non-deterministic because they "introduce noise into their computations, so the computations are not repeatable."

Pet. 22 (quoting Ex. 1001, 4:11–17, 14:19–64, 17:12–18).

Specifically, Petitioner contends that, for "non-deterministic embodiments," the statistical mean would be the average of "the different outputs produced by the same operation on the same input." *Id.* But for "deterministic digital embodiments" like Dockser, Petitioner contends that the statistical mean is "the same as the numerical value of the first output signal for any individual execution of the first operation on each specific input, because that output is always the same for any specific input." *Id.* at 23. Consequently, Petitioner contends, repeatedly executing a multiplication operation using Dockser's floating-point multiplier "on the same input (*i.e.*, pair of operands) with the same precision level yields the same result for every execution; therefore, the statistical mean of the outputs is the same as the output for any single execution." *Id.*

With respect to the "exact mathematical calculation" portion of the imprecision limitation, Petitioner contends that, because Dockser performs a reduced-precision multiplication, the result of the operation differs from what would be the exact mathematical result of the operation, namely "the (>32-bit) product that would result if the pair of input 32-bit operands were multiplied without reducing precision." *Id.* at 24 (emphasis omitted).

With respect to the relative error (Y) for fraction of valid inputs (X) portion of the imprecision limitation, Petitioner contends that Dockser teaches “for at least $X=5\%$ of the possible valid inputs to the first operation,” the statistical mean of the results of executing the first operation “differs by at least $Y=0.05\%$ ” from the result of the exact mathematical calculation. *Id.* at 24–38. According to Petitioner, the “possible valid inputs” in Dockser are “the set of possible normal IEEE-754 32-bit single-format numbers forming pairs of operands in input signals to the execution unit that can be multiplied together to produce an output representing a numerical value (rather than, e.g., an overflow/underflow exception).” *Id.* at 25. Petitioner contends that Dockser’s FPP operates at a precision level meeting the claimed X and Y percentages for such input pairs, pointing to Dockser’s description of retaining only some of the bits of a mantissa fraction (e.g., the 9 most-significant bits of a 23-bit fraction, as shown in Figure 2, above) and dropping the remaining “excess bits” (e.g., the other 14 bits). *Id.* at 25–26. Petitioner contends that Dockser teaches dropping the excess bits to reduce precision by either (1) “removing power from storage elements in the FPP’s registers that correspond to the excess (dropped) mantissa bits,” or (2) “removing power from elements within the multiplier logic that computes the product of the operand mantissas.” *Id.* at 26. Petitioner contends that both techniques, used either individually or in combination, teach the recited imprecision. *Id.*; *see also id.* at 27–39. We determine that Petitioner has made a sufficient showing on the current record with respect to the first precision-reducing technique, and need not evaluate Petitioner’s alternative arguments at this time.

Relying on the first technique and Dockser’s example of retaining 9 mantissa bits and dropping 14 mantissa bits, Petitioner contends that an ordinarily skilled artisan “would have understood the output of unpowered storage elements would be tied to zero voltage (*e.g.*, ground), making those 14 ‘excess’ bits zeroes.” *Id.* at 29. Multiplying two operands with excess bits dropped, Petitioner contends, results in an output with reduced precision from the exact product, where the amount of error depends on the number of mantissa bits dropped. *Id.* at 30–31. Petitioner contends that an ordinarily skilled artisan would have understood, by “straightforward math” described in Appendix I.A to the Petition, that “the relative error (the claimed ‘Y’ percentage) of any floating-point number output from Dockser’s reduced-precision multiplication is the same as the relative error of its mantissa, independent of its exponent and sign.” *Id.* at 31.

Petitioner then provides a detailed explanation as to why a person of ordinary skill in the art would have understood Dockser as teaching the recited X and Y percentages of claim 1. *Id.* at 31–32. First, Petitioner contends that “[g]iven the massive number of possible inputs to Dockser’s FPP (including over 70 trillion possible pairs of normal IEEE-754 single-format mantissas), [an ordinarily skilled artisan] would have performed Dockser’s FPP operation in software to determine the fraction X of all possible valid inputs that produce at least the claimed relative error Y when a given number of mantissa bits are dropped.” *Id.* at 32. Petitioner contends that Mr. Goodin wrote such a program to perform reduced-precision multiplication retaining 9 mantissa bits and dropping the 14 excess bits as in Dockser that tested all possible valid mantissa pairs and “produce[d] at least

Y=0.05% relative error for 92.1% of possible valid inputs (greater than X=5%).” *Id.* at 32–33, 66–68 (citing Ex. 1003 ¶¶ 284, 414–428).

Second, Petitioner contends that an ordinarily skilled artisan “would also have understood algebraically that Dockser’s register bit-dropping technique meets [the limitation], by examining the absolute *minimum* relative error produced by zeroing certain mantissa bit positions.” *Id.* at 33. Petitioner contends that over 12% of possible input operands “have a zero as their most-significant (leftmost) mantissa fraction bit and ones as their eighth and ninth fraction bits,” such that retaining only 9 mantissa bits of operands would result in “every input in that 12% produc[ing] at *minimum* 0.097% relative error.” *Id.* at 33, 68–72 (citing Ex. 1003 ¶¶ 286, 429–448).

Claim 1 recites, in a final limitation, “at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.” Ex. 1001, 30:5–6. Petitioner contends that Dockser’s main processor meets the computing device limitation because the processor writes subprecision select bits to the FPP’s control register. Pet. 38. Petitioner contends that, this way, Dockser’s main processor is adapted to control the FPP’s operation by specifying its precision level. *Id.* (citing Ex. 1007 ¶¶ 15, 18, 25, 35; Ex. 1003 ¶¶ 303–304).

Petitioner’s contentions about claim 1 are supported by the testimony of Mr. Goodin. *See id.* at 9–38; Ex. 1003 ¶¶ 184–304.

b) Patent Owner’s Response

Patent Owner makes two arguments in its Preliminary Response disputing Petitioner’s contentions about claim 1. Prelim. Resp. 8–24. We address each argument separately below.

(1) LPHDR Execution Unit

First, Patent Owner argues that Dockser’s FPP is not a “low precision high dynamic range (LPHDR) execution unit” because it is also “capable of operating at full precision.” Prelim. Resp. 11–12. According to Patent Owner, “Dockser discloses a 32-bit FPP that includes all of the circuitry needed for full precision arithmetic on data in IEEE 32-bit format, and also having additional circuitry allowing for selectable subprecisions.” *Id.* at 11. As support for this reading of Dockser, Patent Owner points to statements in Dockser that precision “may” be reduced and that certain applications require “greater precision.” *Id.* at 11–12 (citing Ex. 1007 ¶¶ 3, 14, 26, 28). Patent Owner also argues that Petitioner incorrectly focuses on whether the FPP is “capable of performing a few operations that are ‘low precision’” in a “9-bit subprecision mode,” rather than “whether the *execution unit itself* can be fairly characterized as being ‘low precision.’” *Id.* at 13. Patent Owner argues that Dockser’s FPP is “a full precision processor that—even when performing operations in reduced precision mode—includes all of the circuitry and capability needed to perform full precision operations.” *Id.*

We have considered Patent Owner’s arguments, but disagree on the current record. We understand Patent Owner’s position as, to be an LPHDR execution unit, the execution unit must be capable of low-precision operations and nothing else. We are not persuaded, on this record, that the claim language is so limiting. Claim 1 recites a “device comprising: at least one first low precision high dynamic range (LPHDR) execution unit.” Ex. 1001, 29:54–56. The only limitations on the execution unit recited in the claim are that the execution unit be “low precision,” “high dynamic range,” and “adapted to execute a first operation” meeting certain criteria specified

in the imprecision limitation (i.e., a minimum relative error Y for a minimum fraction X of possible valid inputs in a specified dynamic range). *Id.* at 29:55–30:4. As Petitioner points out, the claim does not recite any structural characteristics of the execution unit and does not include any negative limitation precluding the execution unit from performing other types of operations. *See* Pet. 31 (“[c]laim 1 recites no *structural* characteristics of the execution unit”); Reply 1–2 (“there is no negative limitation in the claims”). The recitation of “a first operation” in claim 1 further requires only one or more first operations (that meet the low precision criteria specified in the claim); it does not say that “every” operation must be low precision or exclude other capabilities for other operations. *See KCJ Corp. v. Kinetic Concepts, Inc.*, 223 F.3d 1351, 1356 (Fed. Cir. 2000) (“This court has repeatedly emphasized that an indefinite article ‘a’ or ‘an’ in patent parlance carries the meaning of ‘one or more’ in open-ended claims containing the transitional phrase ‘comprising.’”). On this record, we see no reason why the claim precludes the execution unit from having additional circuitry for performing other types of operations, as long as the execution unit is capable of performing the recited first operation and meeting the criteria set forth in the imprecision limitation.

Petitioner has provided sufficient evidence, at this stage of the proceeding, that Dockser performs operations at the precision level specified in the imprecision limitation. Dockser’s FPP “perform[s] certain mathematical operations,” such as “multiplication,” at “reduced precision.” Ex. 1007 ¶¶ 1, 30–32. Specifically, Petitioner expressly identifies “reduced-precision multiplication” as the “first operation” in Dockser. Pet. 15–20; *see also* Prelim. Resp. 10 (acknowledging that Dockser’s FPP performs “a

reduced-precision operation”), 11 (acknowledging that the FPP has “additional circuitry [that] allows the FPP to perform some operations with selectively reduced precision”). And Petitioner calculates the relative error for multiplication of all possible input operands when the operands have 9 retained mantissa bits and 14 dropped mantissa bits.⁴ Pet. 24, 29–33. That is consistent with Dockser, which discloses a specific example of multiplying 23-bit operands having 9 retained mantissa bits and 14 dropped mantissa bits. Ex. 1007 ¶¶ 26, 29.

Importantly, as Patent Owner acknowledges, Dockser’s example of 9-bit precision multiplication is a distinct operation from multiplication at other levels of precision. *See* Prelim. Resp. 20–21 (arguing that Dockser can “perform an addition operation $a+b$ in full precision mode, 22-bit mode, 21-bit mode, etc.” and that “[e]ach of these operations is distinct, and generally produces different outputs, with some being more precise than others”). Dockser describes exactly how such multiplication is performed, and the functionality of the FPP differs depending on which level of precision multiplication is selected. *See, e.g.*, Ex. 1007 ¶¶ 30–34, Fig. 3B; Pet. 23–32. On this record, we find the fact that Dockser is capable of performing other, different operations (e.g., 23-bit full precision) does not detract from Dockser’s disclosure of a specific example meeting the “low precision” requirements of the claim.

Patent Owner further argues that Petitioner fails to show that it would have been obvious to modify Dockser’s FPP to be an LPHDR execution

⁴ Accordingly, we do not agree with Patent Owner that Petitioner improperly made a new argument in its Reply that “multiplying two input values in a way that reduces precision down to 9 mantissa bits” constitutes a “first operation.” Sur-Reply 2.

unit. Prelim. Resp. 14–17. Patent Owner contends that Dockser teaches away from such a unit because

rather than committing itself to performing low precision operations on a very high percentage of all possible valid inputs (in order to shrink the size of the execution unit), Dockser specifically teaches that its selectable precision (which means supporting full precision and programmability, both of which increase the size of the execution unit) is the key feature of the Dockser FPP.

Id. at 15 (citing Ex. 1007 ¶ 3). According to Patent Owner, Dockser’s FPP “is more complex than conventional 32-bit microprocessors, not less,” so an ordinarily skilled artisan would not have been motivated to “reverse course” and make it only capable of low precision operations. *Id.* at 17 (emphasis omitted). As explained above, however, we disagree on this record with Patent Owner’s position that an LPHDR execution unit must be incapable of anything other than low precision operations, and thus do not agree that a modification to Dockser in that respect would have been necessary. We instead find Petitioner’s analysis as to why an ordinarily skilled artisan would have understood Dockser’s FPP to be an LPHDR execution unit, supported by the testimony of Mr. Goodin, sufficient at this early stage.

(2) Imprecision Limitation

Second, Patent Owner argues that Dockser does not teach or suggest the imprecision limitation because Petitioner fails to show that Dockser’s output over all “possible valid inputs” meets the limitation. Prelim. Resp. 18–24 (emphasis omitted). According to Patent Owner, “[t]he range of possible valid inputs for Dockser” includes all possible IEEE 32-bit values “across the entire range of the possible subprecision select bits,” not just operands with 9 retained mantissa bits and 14 dropped mantissa bits as Petitioner discussed in the Petition. *Id.* at 20. Patent Owner argues that

Mr. Goodin’s software program addresses “only a small subset of Dockser’s possible valid inputs” and fails to account for “whether the imprecision limitation is met for subprecision select bits corresponding to any implementation where fewer than 14 bits are dropped.” *Id.* In Patent Owner’s view, the analysis of Dockser needs to take into account operations at “full precision mode, 22-bit mode, 21-bit mode, etc.” *Id.* at 20–21.

We have considered Patent Owner’s arguments, but disagree on the current record. The imprecision limitation recites that “for at least $X=5\%$ of the possible valid inputs to *the first operation*, the statistical mean, over repeated execution of *the first operation* on each specific input from the at least $X\%$ of the possible valid inputs to *the first operation*,” “executing *the first operation* on that input differs by at least $Y=0.05\%$ from the result of an exact mathematical calculation of *the first operation* on the numerical values of that same input.” Ex. 1001, 29:62–30:4 (emphases added). The relevant inquiry thus looks at all possible valid inputs to the first operation—not all possible valid inputs to the LPHDR execution unit overall, or all possible valid inputs to other operations. Again, in its mapping of Dockser to claim 1, Petitioner expressly identifies “reduced-precision multiplication” as “the first operation.” Pet. 15–20. Dockser discloses a specific example of multiplying operands having 9 retained mantissa bits and 14 dropped mantissa bits. Ex. 1007 ¶¶ 26, 30–34, Fig. 3B. The relevant possible valid inputs for assessing the imprecision limitation, therefore, are the inputs to that operation. On this record, we see no reason why the analysis of Dockser also needs to account for other operations, such as “full precision” multiplication, which Patent Owner acknowledges are “distinct” operations.

See Prelim. Resp. 21 (“Each of these operations is distinct, and generally produces different outputs, with some being more precise than others.”).

Patent Owner also challenges Mr. Goodin’s analysis as “impermissibly applying hindsight and using the claims as a roadmap” because “he opines that [an ordinarily skilled artisan] would operate Dockser with the *express goal* of dropping enough bits from the mantissa to meet the imprecision limitation.” *Id.* at 16 (citing Ex. 1003 ¶¶ 280–281). We disagree based on the current record. The precision level of 9 retained mantissa bits and 14 dropped mantissa bits that Mr. Goodin analyzes was not chosen with the goal of meeting the imprecision limitation, but rather is the specific example described in Dockser. *See* Ex. 1003 ¶¶ 255, 269–270, 274–276. Mr. Goodin describes the calculations performed by his software program and explains why he believes that an ordinarily skilled artisan would have understood Dockser’s description of multiplication at the 9-bit precision level teaches the imprecision limitation. *See id.* ¶¶ 281–286, 414–448. Further, Mr. Goodin cites extensively to language in Dockser describing reduced-precision multiplication in support of his opinions. *See id.* ¶¶ 248–259. At this stage of the proceeding, Patent Owner has not presented any evidence indicating that Mr. Goodin’s calculations for Dockser’s 9-bit precision level are factually incorrect. We have reviewed that supporting testimony and are persuaded, based on the current record, that an ordinarily skilled artisan would have understood Dockser to teach the imprecision limitation of claim 1.

c) Claims 2 and 16

We have reviewed Petitioner’s contentions regarding claims 2 and 16 and are persuaded that Petitioner has made a sufficient showing at this stage

for those claims as well. *See* Pet. 38–40. Petitioner explains how each limitation of these claims is taught or rendered obvious by the disclosure of Dockser. *Id.* For example, claim 2 recites that “the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine.” Ex. 1001, 30:7–11. As Petitioner notes, Dockser discloses a general purpose processor. Ex. 1007 ¶ 35; *see* Pet. 39.

Petitioner’s contentions are also supported by the testimony of Mr. Goodin and are persuasive based on the current record. Pet. 38–40; Ex. 1003 ¶¶ 305–310. Patent Owner does not argue the challenged claims separately, only disputing Petitioner’s arguments regarding independent claim 1. *See* Prelim. Resp. 8–34. We disagree with Patent Owner’s arguments based on the current record for the reasons explained above. *See supra* § III.D.1.b.

d) Summary

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 1, 2, and 16 are unpatentable over Dockser.

2. Asserted Obviousness Over Dockser and Tong

Petitioner contends that claims 1, 2, 16, and 33 of the ’156 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over Dockser and Tong. Pet. 40–47. Patent Owner opposes. Prelim. Resp. 24–25. On the present record, for the reasons detailed below, we determine that Petitioner has shown sufficiently for institution that Dockser combined with Tong teaches

or suggests the limitations of the challenged claims and provides sufficient reasoning for combining the references in the manner asserted.

a) Claim 1

Petitioner asserts that “Tong, like Dockser,” “confirms that the number of mantissa bits used in a high-dynamic-range floating-point execution unit was a well-known result-effective variable impacting power consumption and precision.”⁵ Pet. 41–42 (citing Ex. 1008, 273–278; Ex. 1003 ¶¶ 312–313).

Petitioner relies on Tong’s Figure 6, reproduced below with Petitioner’s annotations. *Id.* at 42.




Fig. 6. Program accuracy across various mantissa bitwidths.

Tong’s Figure 6 is a line graph showing program accuracy, from 0% to 100%, on the vertical axis and mantissa bitwidth, from 1 to 23, on the

⁵ Petitioner presents sufficient evidence on the current record to establish a reasonable likelihood that Tong is a prior art printed publication under 35 U.S.C. § 102(b). See Pet. 40–41 (citing Ex. 1025 ¶¶ 8–11; Ex. 1026, 27; Ex. 1027, 27).

horizontal axis. Ex. 1008, 279. The figure shows this data for five programs: ALVINN, Bench22, Fast DCT, PCASYS, and Sphinx. *Id.* The programs implement different signal-processing tasks. *Id.* at 278 (Table IV). ALVINN, for example, is a neural network trainer that uses backpropagation. *Id.* And Sphinx is a speech-recognition program. *Id.*

For the ALVINN and Sphinx line plots in Tong’s Figure 6, Petitioner adds a dashed red line extending vertically through the data points with a mantissa bitwidth of 5. Pet. 42. Petitioner asserts that, for these programs, Tong teaches that “the accuracy does not change significantly with as few as 5 mantissa [fraction] bits.” *Id.* (quoting Ex. 1008, 278–279, 273, 282). In Petitioner’s view, Tong omits unnecessary bits to reduce waste and power consumption. *Id.* at 43. Petitioner characterizes Tong as “[h]aving empirically determined the minimum number of mantissa bits necessary to maintain acceptable accuracy of particular applications.” *Id.* (citing Ex. 1008, 273, 274, 279, 284).

Petitioner contends that “Tong’s teaching that a precision level retaining 5 mantissa fraction bits is sufficient in some applications (including ALVINN and Sphinx[]) would have motivated [an ordinarily skilled artisan] to configure Dockser’s FPP . . . to operate at a selected precision level retaining as few as 5 mantissa fraction bits.” *Id.* at 43–44. According to Petitioner, an ordinarily skilled artisan would have done so “to conserve power when running those applications, or others empirically determined (using Tong’s techniques) to not require greater precision using Tong’s techniques.” *Id.* at 44 (citing Ex. 1003 ¶ 323). Petitioner also contends that “[d]etermining the optimum range of imprecision to achieve the best power reduction without sacrificing accuracy for a particular

application was a matter of routine optimization of a result-effective variable.” *Id.* at 45 (citing Ex. 1003 ¶ 327).

In response, Patent Owner argues that Dockser is deficient for the same reasons discussed in connection with the challenge based on Dockser alone and that Tong does not remedy those deficiencies. Prelim. Resp. 24–25. In Patent Owner’s view, Petitioner’s Dockser-Tong analysis, like the analysis of Dockser alone, accounts for “only a cherry-picked subset of the ‘possible valid inputs.’” *Id.* at 24 (citing Pet. 44–45). Patent Owner argues that Petitioner’s analysis ignores the remaining possible valid inputs across all subprecision selections. *Id.* at 24–25.

Patent Owner’s argument about all subprecision selections, however, does not squarely address Petitioner’s obviousness rationale—i.e., that one of ordinary skill in the art would have been motivated to configure Dockser’s FPP at a “selected precision level” according to Tong. *See* Pet. 43–45. Specifically, Petitioner’s rationale proposes using a selected precision level “retaining as few as 5 mantissa fraction bits.” *Id.* Instead of addressing the analysis of Dockser at this selected precision level, Patent Owner’s arguments focus on Dockser alone. *See* Prelim. Resp. 24–25.

Petitioner’s relative-error analysis of the Dockser-Tong combination, though, is different from the one in the ground based on Dockser alone. *See* Pet. 44. In particular, Petitioner used a software program to determine the relative error when retaining 5 mantissa fraction bits—i.e., the number of bits that Dockser would use under Petitioner’s proposed combination with Tong. *Id.* (citing Ex. 1003 ¶ 324). Also, the algebraic analysis cited by Petitioner uses 5 mantissa fraction bits. *Id.* (citing Ex. 1003 ¶ 324). Because Petitioner’s rationale relies on modifying Dockser based on Tong to have a

particular precision level, we are not persuaded at this stage that Petitioner’s relative-error analysis is deficient for not considering all possible subprecision selections. *See* Prelim. Resp. 24–25; *supra* § III.D.1.b.2. (explaining why Petitioner’s asserted ground based on Dockser alone is not deficient for failing to address all possible IEEE 32-bit values, rather than just operands with 9 retained mantissa bits and 14 dropped mantissa bits).

b) Claims 2, 16, and 33

We have reviewed Petitioner’s contentions regarding claims 2, 16, and 33, and are persuaded that Petitioner has made a sufficient showing at this stage for those claims as well. *See* Pet. 43–47. Petitioner explains how each limitation of these claims is taught or rendered obvious by the disclosure of Dockser or the combination of Dockser and Tong. *Id.* Petitioner’s contentions are also supported by the testimony of Mr. Goodin and are persuasive based on the current record. Pet. 38–40; Ex. 1003 ¶¶ 321–335. Patent Owner does not argue the challenged claims separately, only disputing Petitioner’s arguments regarding independent claim 1. *See* Prelim. Resp. 8–34. We disagree with Patent Owner’s arguments based on the current record for the reasons explained above. *See supra* § III.D.1.b.

Patent Owner makes an additional argument regarding independent claim 33, which recites a “device comprising a computer processor and a computer-readable memory storing computer program instructions, wherein the computer program instructions are executable by the processor to emulate a second device comprising” an LPHDR execution unit having the same properties as recited in claim 1. Ex. 1001, 32:20–42. Patent Owner argues that an ordinarily skilled artisan “would not have been motivated to incorporate Tong’s ‘emulation’ into the systems of Dockser” because “Tong

uses software emulation only for the purposes of *investigating* the behavior of *physical* devices.” Prelim. Resp. 25 (citing Ex. 1008, 273, 278). We disagree based on the current record. Petitioner’s position is that “Tong’s teaching to ‘emulate[] in software different bitwidth FP units’ ‘to determine application accuracy’ . . . would have motivated [an ordinarily skilled artisan] to *emulate the Dockser/Tong device . . . in software* to assess the accuracy of applications running on the device at selected precision levels.” Pet. 46 (quoting Ex. 1008, 278; citing Ex. 1003 ¶¶ 328–329) (emphasis added). Emulating the device of Dockser (modified based on Tong) in software appears, on this record, to be consistent with Tong, which describes the emulation of a floating-point (FP) device to determine “the relationship between program accuracy and number of bits in FP representation.” Ex. 1008, 278. Based on the current record, we do not agree that Tong’s disclosures regarding FP hardware would have dissuaded an ordinarily skilled artisan from making the combination that Petitioner proposes.

c) Summary

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 1, 2, 16, and 33 are unpatentable over Dockser and Tong.

3. Asserted Obviousness Over Dockser and MacMillan

Petitioner contends that claims 1–8 and 16 of the ’156 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over Dockser and MacMillan. Pet. 47–55. Patent Owner opposes. Prelim. Resp. 25–27. On the present record, for the reasons detailed below, we determine that Petitioner has shown sufficiently for institution that Dockser combined with MacMillan teaches or suggests the limitations of the challenged claims and

provides sufficient reasoning for combining the references in the manner asserted.

a) Claim 1

In this asserted ground, Petitioner relies on MacMillan for its teachings about multiple floating-point execution units and contends that it would have been obvious to implement a device with multiple Dockser FPPs operating in parallel. Pet. 47–50.

Patent Owner argues that MacMillan does not remedy Dockser’s deficiencies. Prelim. Resp. 25–27. Apart from this argument, Patent Owner does not present arguments specifically analyzing MacMillan as it is used in Petitioner’s challenge to claim 1. *Id.* Rather, Patent Owner refers to the arguments regarding the challenge based on Dockser alone. *Id.* at 25–26. Thus, for the reasons discussed in connection with the asserted ground based on Dockser alone, and because we conclude that Petitioner has made a sufficient showing of obviousness based on the combination with MacMillan, we determine on this record that Petitioner has shown a reasonable likelihood of prevailing on its contention that claim 1 is unpatentable over Dockser and MacMillan.

b) Claim 3

Claim 3 depends from claim 1 and recites “wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.” Ex. 1001, 30:13–17. Petitioner contends that MacMillan teaches 256 processing elements (PEs), and the Dockser-MacMillan combination would have “a single Host CPU and at least one

FPP in each PE (of which there are at least 256).” Pet. 52 (citing Ex. 1003 ¶¶ 361–362). Thus, in Petitioner’s proposed combined device based on Dockser and MacMillan, the number of LPHDR execution units (Dockser FPPs) “exceeds by over 100 its number (one) of traditional-precision execution units (the single Host CPU floating-point unit).” *Id.* at 52–53 (citing Ex. 1003 ¶ 366).

Patent Owner argues that “Dockser and MacMillan cannot possibly disclose the limitation of claim 3, even if one were to assume the Dockser FPP to be an LPHDR execution unit.” Prelim. Resp. 27. Patent Owner argues that, because (1) Petitioner asserts that MacMillan’s processor is the only execution unit (EU) for purposes of the claims, and (2) Dockser’s FPPs are designed to perform multiplication on 32-bit numbers, “the number of claim 3 32-bit EUs in the Dockser/MacMillan combination must *always* be greater than or equal to the number of ‘LPHDR’ Dockser FPPs.” *Id.*

But, under Petitioner’s theory, the identified 32-bit EUs in claim 3 (i.e., “execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide”) are different from what Petitioner identifies as the LPHDR execution units (i.e., a Dockser FPP in each of the 256 (or more) PEs of MacMillan). *See* Pet. 51–52. Thus, under Petitioner’s rationale, Dockser’s FPPs do not count toward the number of 32-bit EUs in the Dockser-MacMillan combination. *See id.* Specifically, Petitioner has shown, on this preliminary record, that the claimed “low precision high-dynamic range (LPHDR) execution unit” at least encompasses Dockser’s FPP. *See supra* § III.D.1.a. Petitioner contends that the ’156 patent describes the LPHDR execution units as “sometimes” producing results that are different from the correct result, in

contrast to 32-bit EUs, which the '156 patent describes as “‘traditional precision’ execution units that do not ‘sometimes’ produce results different from the correct traditional-precision result.” Pet. 52. Thus, if the proposed combination has 256 (or more) FPPs (LPHDR execution units) and one host CPU (the 32-bit EUs), the number of LPHDR execution units would be greater than the number of 32-bit EUs by the amount recited in claim 3.

Petitioner supports its contentions with respect to claim 3 with citations to the '156 patent. We preliminarily agree with Petitioner that the cited parts of the patent distinguish between the two sets of units. For example, the '156 patent describes the 32-bit arithmetic elements as “‘traditional’ precision: “‘arithmetic elements . . . designed to perform . . . floating point arithmetic with a word length of 32 or more bits’ are ‘designed to perform . . . arithmetic of traditional precision.’” Ex. 1001, 27:52–63). As for the LPHDR units, the paragraph cited by Petitioner explains how the precision may vary across implementations. *See id.* at 26:61–27:17. We preliminarily agree with Petitioner that all these embodiments, unlike the 32-bit arithmetic elements, are described as producing results that are “‘sometimes” or “all of the time” not closer than a certain amount to the correct result. *See* Pet. 52–53 (citing Ex. 1001, 26:39–49). Thus, Petitioner’s distinction between the 256 (or more) FPPs and one host CPU in the proposed combination is consistent with the language of the challenged claims and adequately supported on this preliminary record.

c) Claims 2, 4–8, and 16

We have reviewed Petitioner’s contentions regarding claims 2, 4–8, and 16 and are persuaded that Petitioner has made a sufficient showing at this stage for those claims as well. Pet. 50–51, 53–55. Petitioner explains

how each limitation of the dependent claims is taught or rendered obvious by the disclosure of Dockser and MacMillan. *Id.* Petitioner's contentions are supported by the Goodin Declaration and are persuasive based on the current record. Patent Owner does not argue claims 2, 4–8, or 16 separately. *See* Prelim. Resp. 24–27. We disagree with Patent Owner's arguments based on claim 1 for the reasons explained above. *See supra* § III.D.1.b.

d) Summary

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 1–8 and 16 are unpatentable over Dockser and MacMillan.

4. Asserted Obviousness Over Dockser, Tong, and MacMillan

Petitioner contends that claims 1–8, 16, and 33 of the '156 patent are unpatentable under 35 U.S.C. § 103(a) as obvious over Dockser, Tong, and MacMillan. Pet. 56–60. Patent Owner opposes. Prelim. Resp. 28–32. On the present record, for the reasons detailed below, we determine that Petitioner has shown sufficiently for institution that Dockser combined with Tong and MacMillan teaches or suggests the limitations of the challenged claims and provides sufficient reasoning for combining the references in the manner asserted.

a) Claims 1–8, 16, and 33

Petitioner contends that it would have been obvious to operate the FPPs in the Dockser-MacMillan combination at Tong's precision levels. Pet. 56–57 (citing Ex. 1008, 278, Table IV). Petitioner contends that using Dockser's FPP with Tong's precision levels in MacMillan's multiple PEs would have achieved "supercomputer performance" while conserving power. *Id.* (citing Ex. 1003 ¶ 387).

In response, Patent Owner argues that “Petitioner relies on the same flawed reasoning as discussed [in the Preliminary Response] with respect to Dockser and Tong,” but does not present separate arguments directed to the combination of Dockser, Tong, and MacMillan. Prelim. Resp. 28. We disagree with Patent Owner’s arguments based on the current record for the reasons explained above. *See supra* §§ III.D.1, III.D.2. On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 1–8, 16, and 33 are unpatentable over Dockser, Tong, and MacMillan.

b) Petitioner’s Alternative Interpretation of Claim 3

Claim 3 depends from claim 1 and recites “wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device *adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.*” Ex. 1001, 30:13–17 (emphasis added). Petitioner provides an alternative interpretation of the “adapted-to” clause. Pet. 57–60. In particular, Petitioner argues that a unit could meet the adapted-to clause if it is capable of 32-bit multiplication in *some* configurations. *Id.* at 58–60. Under this interpretation, Dockser’s FPP would not be such a unit because it has registers and multiplier with less than 32 bits, and thus is not capable of 32-bit multiplication in the proposed combination—rather, only the host CPU floating-point unit would meet the adapted-to clause. *Id.* (citing Ex. 1003 ¶¶ 400–401). According to Petitioner, an ordinarily skilled artisan “would have been motivated to customize Dockser’s FPPs in MacMillan’s PEs to only operate at precision

levels lower than full FP 32-bit operations, in view of Tong’s teachings that ‘the fine precision of the 23-bit mantissa is not essential.’” *Id.* at 58.

Patent Owner argues that Petitioner does not provide a specific level of precision for its combination. Prelim. Resp. 31 (citing Pet. 58–59). Patent Owner argues that “Tong lists precisions ranging between 5 and 11 bits, but Petitioner does not state which, if any of those, [an ordinarily skilled artisan] would use.” *Id.* at 32. In addition, Patent Owner argues that Petitioner does not show that, for example, one bit lower than full 32-bit operation would meet the claimed imprecision limitation. *Id.* at 31–32.

We have considered Patent Owner’s arguments, but disagree on the current record. In its “alternative interpretation of claim 3,” Petitioner states that the “selected precision levels are unchanged from [the other three asserted grounds].” Pet. 59. In the asserted ground based on Dockser and Tong, for example, Petitioner states that Tong’s “teaching that a precision level retaining 5 mantissa fraction bits is sufficient in some applications (including ALVINN and Sphinx[]) would have motivated [an ordinarily skilled artisan] to configure Dockser’s FPP . . . to operate at a selected precision level retaining as few as 5 mantissa fraction bits.” *Id.* at 43–44. Petitioner also uses a software program to determine the relative error when retaining 5 mantissa fraction bits—i.e., the number of bits that Dockser would use under Petitioner’s proposed combination with Tong. *Id.* at 44 (citing Ex. 1003 ¶ 324). In at least this way, Petitioner has adequately explained, for institution and on this preliminary record, which of Tong’s precision levels would be used in the proposed combination.

Patent Owner also argues that each reference teaches away from the proposed combination. Prelim. Resp. 30. In particular, Patent Owner argues

that (1) Dockser teaches away because “Dockser is directed entirely to selectable precision; and disparages the use of non-selectable precision units,” (2) Tong teaches away because “Tong devotes much of its discussion to the benefits of variable or selectable precision,” and (3) MacMillan teaches away because “MacMillan is primarily concerned with providing increased performance without increased cost.” *Id.* at 30–31.

At this stage and on this record, we preliminarily determine that the references do not teach away from the claimed devices. The mere disclosure of more than one alternative does not constitute a teaching away. *In re Fulton*, 391 F.3d 1195, 1201 (Fed. Cir. 2004). Specifically, Patent Owner points to a statement that Tong makes about power savings: “This FP bitwidth reduction can deliver a significant power savings through the use of a variable bitwidth FP unit.” Prelim. Resp. 31 (quoting Ex. 1008, 273). Yet Tong’s statement does not discredit, criticize, or otherwise discourage the claimed invention’s approach so as to teach away from it. *See Fulton*, 391 F.3d at 1201. Instead, the statement merely explains the benefits provided by Tong’s alternative approach. Ex. 1008, 273.

Likewise, although Patent Owner argues the combination with MacMillan would increase manufacturing costs, Patent Owner does not provide sufficient evidence to support this position on this record. *See* Prelim. Resp. 31. Even assuming MacMillan warns against the costs, Patent Owner’s argument is unsupported by any specific cost analysis. *Id.* To be sure, Patent Owner may introduce evidence that supports this argument during trial. We also note that the fact “[t]hat a given combination would not be made by businessmen for economic reasons does not mean that persons skilled in the art would not make the combination because of some

technological incompatibility,” for example. *In re Farrenkopf*, 713 F.2d 714, 718 (Fed. Cir. 1983). But, at this stage and on this record, Patent Owner’s argument is unpersuasive.

c) Summary

On this record, we are persuaded that Petitioner has shown a reasonable likelihood of prevailing on its assertion that claims 1–8, 16, and 33 are unpatentable over the combination of Dockser, Tong, and MacMillan.

IV. CONCLUSION

After considering the arguments presented in the parties’ papers and the evidence of record, we determine that Petitioner has demonstrated at least a reasonable likelihood of success in proving that at least one claim of the ’156 patent is unpatentable. Thus, we institute an *inter partes* review of all challenged claims (i.e., claims 1–8, 16, and 33) on all grounds set forth in the Petition. Our determinations at this stage of the proceeding are based on the evidentiary record currently before us. This decision to institute trial is not a final decision as to patentability of any claim for which we have instituted an *inter partes* review. *See TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1068 (Fed. Cir. 2016) (noting that “there is a significant difference between a petitioner’s burden to establish a ‘reasonable likelihood of success’ at institution, and actually proving invalidity by a preponderance of the evidence at trial”). We will base any final decision on the full record developed during trial.

V. ORDER

Accordingly, it is

ORDERED that, pursuant to 35 U.S.C. § 314(a), an *inter partes* review is instituted for claims 1–8, 16, and 33 of the '156 patent on the unpatentability grounds asserted in the Petition; and

FURTHER ORDERED that pursuant to 35 U.S.C. § 314(c) and 37 C.F.R. § 42.4, notice is hereby given of the institution of a trial, which commences on the entry date of this decision.

IPR2021-00165

Patent 9,218,156 B2

FOR PETITIONER:

Elisabeth Hunt

Richard Giunta

Anant Saraswat

WOLF, GREENFIELD & SACKS P.C.

ehunt-ptab@wolfgreenfield.com

rgiunta-ptab@wolfgreenfield.com

asaraswat-ptab@wolfgreenfield.com

FOR PATENT OWNER:

Peter Lambrianakos

Vincent Rubino

Enrique Iturralde

FABRICANT LLP

plambrianakos@fabricantllp.com

vrubino@fabricantllp.com

eiturralde@fabricantllp.com